

# EXHIBIT W

# **EXHIBIT 57**

**EXHIBIT FILED UNDER SEAL**

## UBER\_JCCP\_MDL\_002340895

## Metadata

#Author	jennifera@uber.com	SEMANTIC
#Date Modified	12/08/2022	SEMANTIC
#DateCreated	04/07/2022	SEMANTIC
#Title	SRAD SOFT FILTERING FAQ FOR MTR	SEMANTIC
Account	mike.akamine@uber.com;	SEMANTIC
All Custodians	Akamine, Mike;Hasbun, Andrew;Kansal, Sachin;	SEMANTIC
All Paths	Akamine, Mike: \JCCP_Drive001\JCCP_Drive001_13.zip; Akamine, Mike: \JCCP_Drive001\JCCP_Drive001_13.zip; Hasbun, Andrew: \EDISCO-25937_andrew.hasbun@uber.com_mai@uber.com_meghan@uber.com\EDISCO-25937_andrew.hasbun@uber.com_mai@uber.com_meghan@uber.com_16.zip; Hasbun, Andrew: \EDISCO-25937_andrew.hasbun@uber.com_mai@uber.com_meghan@uber.com\EDISCO-25937_andrew.hasbun@uber.com_mai@uber.com_meghan@uber.com_16.zip; Kansal, Sachin: \EDISCO-25937_sachin.kansal@uber.com\EDISCO-25937_sachin.kansal@uber.com2_28.zip; Kansal, Sachin: \EDISCO-25937_sachin.kansal@uber.com\EDISCO-25937_sachin.kansal@uber.com2_28.zip	SEMANTIC
Application	Microsoft 2007 Word Document	SEMANTIC
Attachment Names	word	SEMANTIC
Begin Family	UBER_JCCP_MDL_002340895	SEMANTIC
Collaborators	lfernandez@uber.com; mingxuan@uber.com; meredith.moore@uber.com; aliorab@uber.com; andrew.hasbun@uber.com; dsvirsky@uber.com; dkolita@uber.com; dharmin@uber.com; anapaulapt@uber.com; gorkem@uber.com; vo@uber.com; bmarchi@uber.com; duvlad@uber.com; gferreira@uber.com	SEMANTIC
Confidentiality	Confidential	SEMANTIC
Date Created	04/07/2022 12:12 am	SEMANTIC
Date Modified	12/08/2022 7:19 pm	SEMANTIC
DocID	19kfxoyBOjxRNTM0eIVCdrCif5EyMwky_Sdxw85zwRY	SEMANTIC
Document Type	Electronic File	SEMANTIC
End Family	UBER_JCCP_MDL_002340907	SEMANTIC
File Path	\JCCP_Drive001\JCCP_Drive001_13.zip	SEMANTIC
File Size	1205739	SEMANTIC
Filename	SRAD SOFT FILTERING FAQ FOR MTR_19kfxoyBOjxRNTM0eIVCdrCif5EyMwky_Sdxw85zwRY.docx	SEMANTIC
GoogleDocumentType	DOCUMENT	SEMANTIC
Hash Value	7c632f9890c095c29c36014277018b57	SEMANTIC
Hidden Content	Yes;	SEMANTIC
ILS All Bates	UBER_JCCP_MDL_002340895;UBER_JCCP_MDL_002340896;UBER_JCCP_MDL_002340897;UBER_JCCP_MDL_002340898;UBER_JCCP_MDL_002340899;UBER_JCCP_MDL_002340900;UBER_JCCP_MDL_002340901;UBER_JCCP_MDL_002340902;UBER_JCCP_MDL_002340903;UBER_JCCP_MDL_002340904;UBER_JCCP_MDL_002340905;UBER_JCCP_MDL_002340906;UBER_JCCP_MDL_002340907	SEMANTIC
ILS Document Date	12/08/2022	SEMANTIC
ILS Prod Date	1/23/2025	SEMANTIC
ILS Prod Vol	JCCP_MDL064	SEMANTIC
Other Custodians	Akamine, Mike;Hasbun, Andrew;Kansal, Sachin;	SEMANTIC
Primary Date	04/07/2022 12:12 am	DOC_TYP E ALIAS
Production Volume	JCCP_MDL064;	SEMANTIC
Redacted	No	SEMANTIC
Sort Date	12/08/2022 7:19 pm	SEMANTIC
SourceHash	007b4457ddf0be59863532c8503a3f16	SEMANTIC

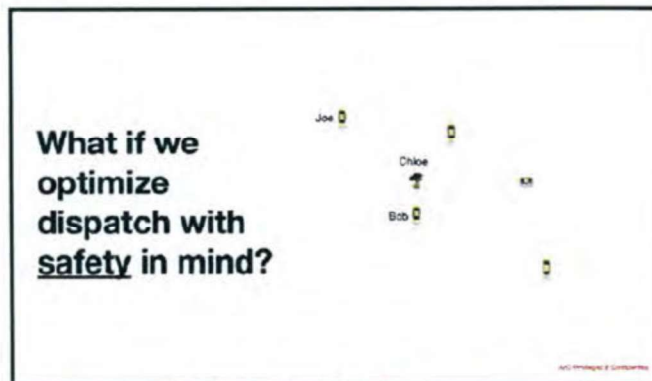


## FAQ for MTR

Resources: [MTR Deck](#)

Presentation: April 21

General questions



- **What do we currently optimize dispatch for?**  
Traditionally, we optimize for ETA and GB (long trips?)
  - **We have been at it for a while. Why has it taken so long? / Why does it take so long to run a safety experiment?**  
Long experiment times are typical for safety experiments that try to measure incident rates for rare events. For S-RAD this gets even more pronounced since:
    - the product is aiming to reduce a subset of all safety incidents, making the target an even rarer event
    - the nature of the soft intervention brings to necessity to reduce [REDACTED]
- For S-RAD:
- [REDACTED]
  - [REDACTED]
  - Covid lockdowns and related trip volatility put the program on ice for a bit.

Commented [1]: @dkolta@uber.com  
@andrew.hasbun@uber.com @anapaulapt@uber.com  
@dsivirsky@uber.com @gorkem@uber.com  
@bmarchi@uber.com @ifernandez@uber.com  
@gferreira@uber.com @duviad@uber.com  
@mingxuan@uber.com

Hi Folks. We are crowdsourcing questions ahead of the MTR. Please take a look at the MTR deck (if you have not seen already) and then share any questions here that pop up from the deck or otherwise. Thanks!

Best,  
Jenny

cc: @dhamin@uber.com



- How did we determine different subcategories?**  
 What we've learned from working with NSVRC, it is essential to make our taxonomy MECE, behaviorally specific, and easily interpreted by agents. That leads us to the most concise number of sub-categories as possible.
- Why are some issues such as Sexual Misconduct and Verbal altercations considered non-serious issues? Shouldn't we have zero tolerance policy for Sexual misconduct?**  
 All safety issues are unacceptable on our platform. However, what we've learned from working with the advocate groups, is that some behavior may be improved via training, while others require a more harsh policy.
- What sort of sexual misconduct training resources do we provide/require for drivers and riders?**

Currently, we have sexual assault and sexual misconduct education for drivers in 51 countries and it's mandatory in about 19 countries. The materials were created by a third party expert, RAINN and other local NGOs in certain markets. Our safety education website, which I've listed next to the question, shows a preview of the lesson and a map on current coverage. It's linked to the source of truth which provide links to all the education content if there is a specific market you're interested in.



- **What is the difference in actions we take when we receive a non-serious versus serious versus critical sasm report?**

For serious/critical incidents, we deactivate.

For non serious incidents, we track the history and have a strikes system.


We have education training, in some markets even mandatory.

Commented [2]: @alirorab@uber.com - can you please provide a link or answer for MTR audience?

- **What are the 5 most severe sexual assault categories?**

1. Non-Consensual Sexual Penetration (NCSP)
2. Attempted NCSP
3. Non-Consensual Kissing - Sexual Body Part
4. Non-Consensual Touching - Sexual Body Part
5. Kissing of a Non-Sexual Body Part

- **What are the top 3 serious sexual assault categories by volume?**



- **What is the difference between Sexual Misconduct & Severe Sexual Misconduct?**  
Severe SM: Masturbation / Indecent exposure




SM: Asking personal questions


- **Why did we decide to focus on serious versus all sasms?**  
Impact on the victim is much more significant  
The volume of non-serious is many fold higher than serious sasm  
Non-serious incidents could crowd out serious incidents in terms of learning for the model
- **Why did we decide not to focus on all IPCs?**  
We have other products such as SDM, Rider verification and DACT policies that help address non-sexual interpersonal conflicts.

SRAD is complicated product; so we wanted to focus the team and model on Sexual assault problem first. Once we have tackled this, we intend to expand the footprint to all IPCs.





- **Have we seen patterns in SRAD trips flagged?**
  - Guest Rides
  - Intoxicated riders
  - Bar trips
  - Yet to be explored further
- **Why are we taking rider & driver features together? Why don't we assess their risks separately?**
  - SRAD is another way of looking at risk at the trip level
  - We believe fundamentally that the risk is situational - there is not something intrinsic of a rider or driver
  - Example: SDM looks at risk of rider (for theft, not SASM). We have DACT (look at riders or drivers based on safety incident history)
  - Dact models in H2 look at user risk.





Commented [3]: @bmarchi@uber.com  
@gorkem@uber.com

- [REDACTED]
  - In conversations with civil rights advocates in the US, [REDACTED]
  - [REDACTED] It's also important to note that removing gender doesn't remove the need to be vigilant for bias. A machine learning model can 'guess' gender if it wants to, which is why we continue to test for gender bias.
- **What is the problem with using gender inferences in the model?**
  - There were policy concerns and one legal concerns with using gender inferences in the model. First, in some states, explicitly using gender raised the scrutiny that the model would receive for disparate impact. Second, advocates felt it was disrespectful to guess someone's gender and use it in a model without getting consent proactively. As one put it – why not just ask people first?
- **What concerns did advocacy groups have?**
  - The main concern is simple discomfort with a strange technology – what right does Uber have to intervene in people's lives in a way that dictated who might or might not get sexually assaulted? And why do this in the background, without giving people a chance to weigh in?
  - Another main concern is racial discrimination. The model should not discriminate against any groups.
  - Secondary concerns were that users should be able to opt-out of gender and racial inference.
  - These variables would be securely stored and never used for marketing purposes.
  - Uber should be transparent about the use of the model and the PII used in it. Recommendations were to include it on the Privacy Notice, What Moves Us, and Download Your Data.
  - There should be external and independent audit of safety and fairness.
- **Which advocacy groups did we talk to?**
  - ReNika Moore, **ACLU Racial Justice Project**
  - Aaron Rieke, **Upturn and FTC**
  - Dariely Rodriguez, **Lawyer's Committee**
  - Gaylynn Burroughs, **Leadership Conference on Civil and Human Rights**
  - Ryan Budish, **Harvard Berkman Klein**
  - Diana Mancera, **Jane Doe (Massachusetts Coalition Against SA and DV)**
  - Sophie Hilgard, **Harvard Data Science Initiative**
  - Chelsea Barabas, **MIT**
  - Vincent Southerland, **Center for Race Inequality and the Law (NYU)**

- Karen Levy, **Cornell University**
- Roel Dobbe, **AI Now Institute**
- Amy Barasch, **Her Justice**
- John Verdi, **Future of Privacy Forum**
- Chad Sniffen, **National Alliance to End Sexual Violence**
- Erica Olsen, **National Network to End Domestic Violence (Safety Net Initiative)**
- Spandana Singh, **New America - Open Technology Institute**
- Liz Woolery, **Center for Democracy and Technology**

- **How is this model different from SDM?**

SDM assesses a risk of a Rider committing a Theft or Robbery on the platform. When SDM risk is higher than the threshold, it blocks rider's access to the trip and we offer a set of rider verification challenges such as DocScan to the rider. If the rider passes the challenge, we grant them access.


SRAD on the other hand assess the risk of a rider-driver pair. [REDACTED]

- **Why is the precision so low? How do we plan to improve the precision?**


This is something related to the inherent unpredictability of the safety incidents. The team's efforts will continue to improve the precision of the model gradually, however the low-precision nature of S-RAD is not expected to change for the foreseeable future.



- [REDACTED]




- **How does SRAD work in conjunction with job boards?**  
Job boards shows trips that are reasonably available. [REDACTED]  
[REDACTED] This is something that we need to deep dive into.



**Commented [4]:** @fernandez@uber.com  
@gferreira@uber.com - do we know the answer to this?  
\_Assigned to fernandez@uber.com\_

**Commented [5]:** No, this is honestly the first time I hear about job boards. Do you have more context about it, so we can deep dive?



- How do you measure fairness?

- For fairness testing, [REDACTED]

- We work with civil rights and other advocacy groups to help us know where to look and what feels fair.

Dan Svirsky

- **How do we know user ethnicity or which drivers identify as having non-binary gender?**

Answer: On ethnicity inferences – we've done innovative work here. We partnered with civil rights NGOs and academics who focus on algorithmic bias to develop best practices for ethnicity inferences. The inference methods we use rely on users' names and locations. No single inference is perfect, so we've found that the best strategy is to use multiple inference approaches and see what they say directionally, taken together. That work is going to be public in June, and our hope is that other tech companies can use it, and that it's a chance for Uber to show leadership.

On gender identity – some states allow drivers to list X as their driver's license gender, and some drivers call in and self-identify as transgender (we then give them specialized customer support). We can look at this population of drivers – about 10,000 in total – to see if SRAD is making things harder on them.


- **Do we look at race or ethnicity? Do we expect these features to change a lot from region to region when we expand?**


Yes – we look at race and ethnicity in the US. We chose not to in Brazil, because we just don't feel comfortable with the amount of data we have. In the US, we can leverage

public work on ethnicity inferences. In Brazil, that doesn't exist to the same degree. That said, the Fairness Research team is working on solutions for international ethnicity testing and we hope to have that ready as SRAD and DACT continue to roll out.

- **Do users provide us with gender information? Do we have their consent to use the data?**

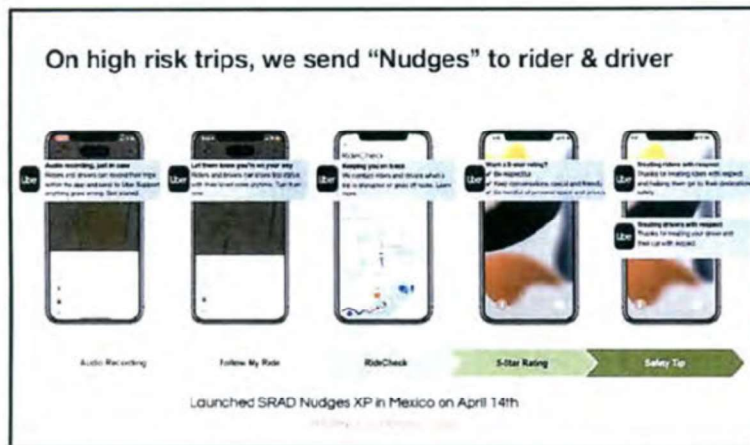
In some cases, we do have user provided gender information. For example, we have driver gender information based on ID documents; but in cases where we don't have user provided gender information, we infer gender based on a user's name.





CONFIDENTIAL

UBER\_JCCP\_MDLE002340005  
#2175712



- What is the long term strategy for nudges?
- What else can we do on risky trips?
- What metrics are you tracking?

Questions from Dara:

1. Does the model pick up features from riders or drivers more?
2. [REDACTED]
  - a. Check that the core data is accurate
3. [REDACTED] @dharmin
4. Will we have the capability to ensure that there is no gender and ethnicity impact?



- a. Yes, we are building Holdout strategy to be able to set up infra to measure all the impacts
    - i. Also bc models deteriorate
    - ii. [REDACTED]
      - 1. We need to balance all the ethics
      - 2. [REDACTED]
  - b. Could we share that we are A/R?
5. Can we talk about this?
- a. We are working on a reactive post
  - b. Malt says its tricky and controversial
    - i. Wants to guard against misinterpretation and criticism
    - ii. Can build it into the safety reports going forward